





National Aeronautics
and Space Administration



System-Level Hardening for Space Systems


Ray Ladbury

NASA Goddard Space Flight Center
Radiation Effects and Analysis Group




Why Look At System Hardening Now?

- Use of commercial parts is a major driver for system hardening
- Use of commercial parts may be justified when
 - Mission has challenging requirements that can't be met with rad-hard parts
 - Commercial product offers important enabling advantage over rad-hard solution







Why Look At System Hardening Now?




- Commercial to rad-hard differential is likely to increase with future generations
- Increased chip complexity means increased testing cost (~3x in 10 years)
- System hardening can decrease test costs and may port across generations








Outline

- 1. Systems and system hardening
 - 1.1. System definition and overview of basic hardening approach
- 2. Threat Evaluation: requirements, technology and inference
 - 2.1. Requirements
 - 2.2. Starting points for system mitigation
 - 2.3 Threat consequences and risk analysis
- 3. Threat Characteristics
 - 3.1 Destructive SEE
 - 3.2 Nondestructive SEE
 - 3.3 Degradation mechanisms
- 4. Mitigation Strategies and Techniques
 - 4.1 Mitigation strategies for destructive SEE
 - 4.2 Mitigation strategies for nondestructive SEE
 - 4.3 Mitigation of degradation mechanisms
- 5. Example: Hardening a memory system for the Solar Dynamics Observatory
- 6. Challenges and conclusions

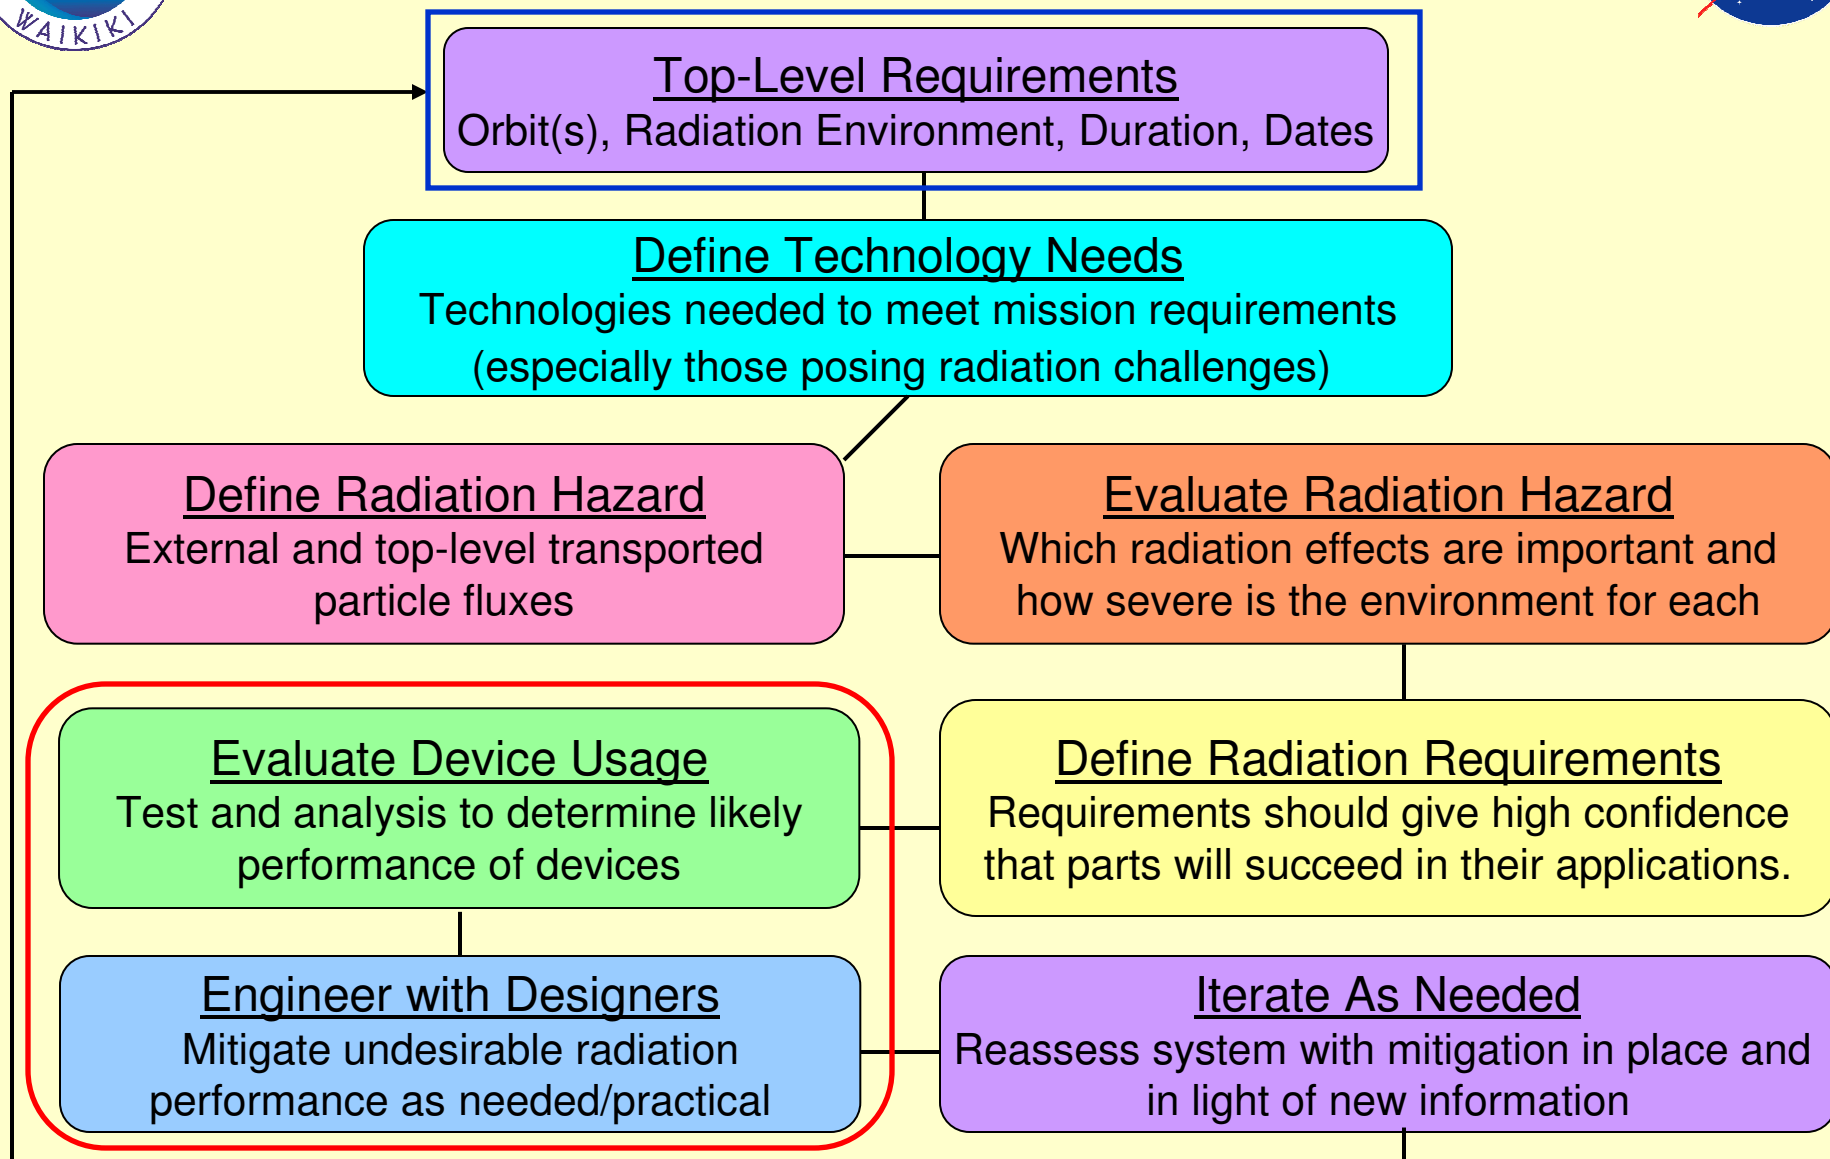



Why System-Level Hardening?





NASA Approach to System-Level Hardening






So, Where Do Requirements Come From?



- Top-level requirements:
 - Mission length, orbit, performance and objectives
 - Usually very general and not restrictive.
- System-level requirements derived from top-level
 - Include survivability, availability, etc.
 - Derived in consultation with system engineers, etc.
- Circuit level requirements:
 - Derived in consultation with design engineers, system engineers reliability and radiation experts, etc.
 - Radiation requirements may be derived from circuit performance requirements and radiation-environment estimates (e.g. TID)
- Verification requirements may be most relevant to radiation tests
 - Need to show that parts will meet their circuit level performance and survivability requirements
 - If part does not meet requirements, we mitigate—effectively easing the requirements on the part.




Hardening Begins with Requirements

- Good Requirements must be:
 1. Clear to all affected parties
 2. Relevant to mission objectives (not “desirements”)
 3. Verifiable by test or analysis (preferably before spacecraft launch)

Top Level


The system shall operate without system-level degradation of capability for 5 years in a geostationary orbit.



Make it relevant to the designer/parts engineer.

Second Level

Parts used in the system shall be immune to destructive SEE.



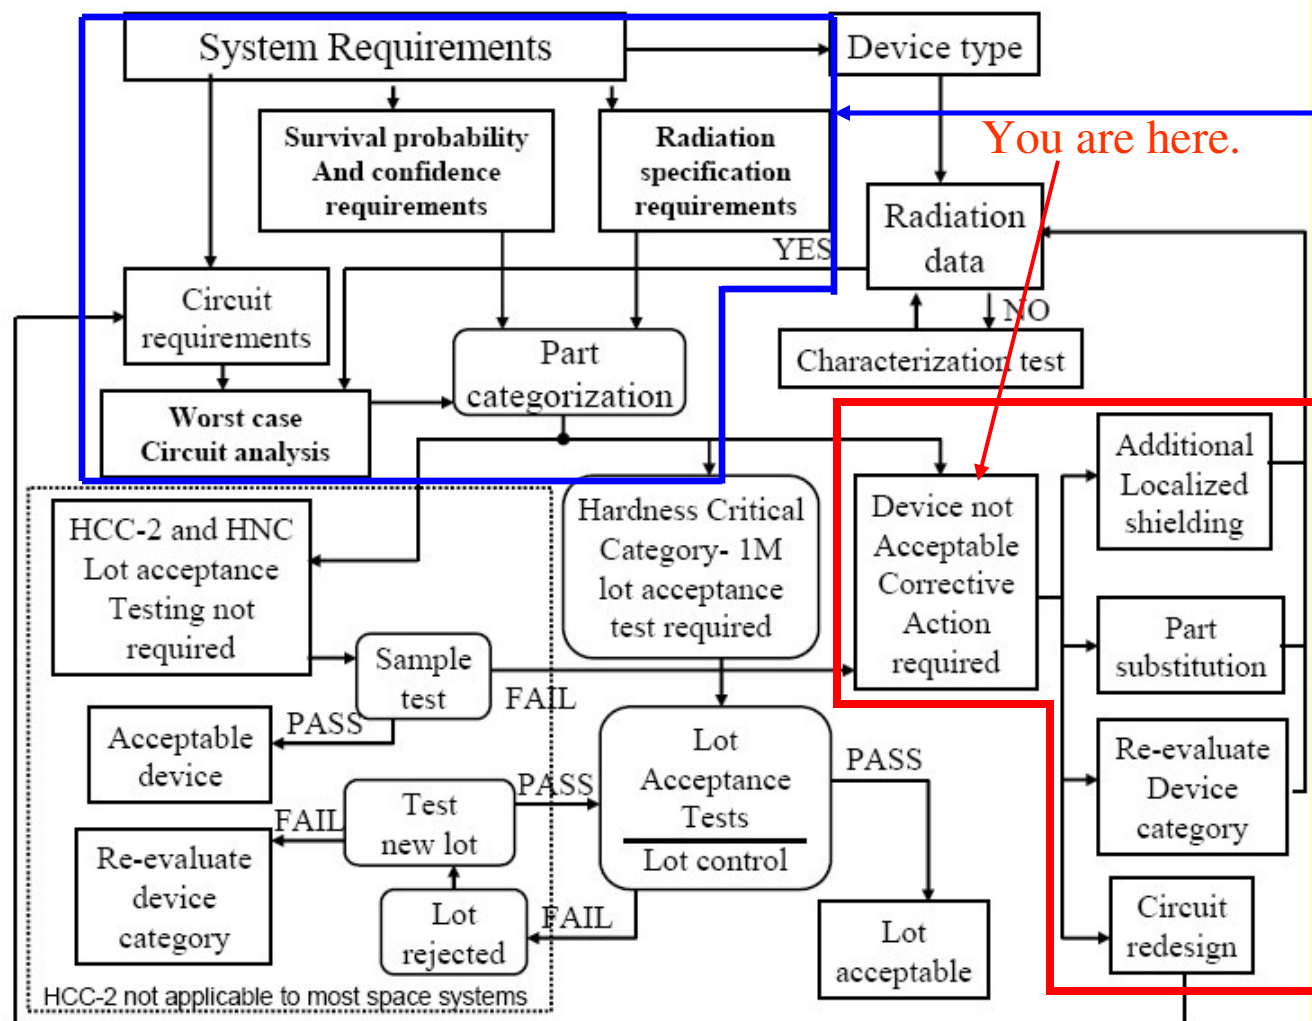

Make it verifiable.

Tertiary Level

Immunity to destructive SEE shall be demonstrated when at least two samples of the part exhibit no destructive SEE after exposure to at least 3×10^7 ions with $LET \geq 60 \text{ MeVcm}^2/\text{mg}$.




Starting Point: TID



Device fails to meet requirements—failing either parametrically or functionally.


“Device not Acceptable” is not a prediction of failure.

Figure 2 from MIL-STD 814




TID Requirements I

- Most TID requirements phrased in terms of design margin, DM
 - $DM = \text{ratio of mean failure dose to application dose}$
 - Usually, minimum acceptable $DM=2$
 - Covers many sins—part variability, environmental uncertainty, etc.
 - Very high reliability applications may demand $DM>2$.





Distribution of increased Ibias for several lots of AD OP484 op amps shows evidence of bimodality.

[23] R. Ladbury and J. Gorelick, TNS 2005




TID Requirements II


- TID requirements sometimes phrased in terms of success probability, P_s and confidence level, CL
 - Requires assuming a form—usually normal or lognormal—for the failure distribution




Problem: Errors on both sample mean and standard deviation depend on the parent standard deviation, σ , which is unknown.



TID Requirements III



- Use one-sided tolerance limits, $K_{tl}(n, P_s, CL)$, for the appropriate test sample size, n , success probability, P_s , and confidence level, CL .
Normal: $D(n, P_s, CL) = m \pm K_{TL}(n, P_s, CL) \times s$
Lognormal: $D_{LN}(n, P_s, CL) = \exp(m_{ln} \pm K_{TL}(n, P_s, CL) \times s_{ln})$
- Parts Characterization Criterion* for TID is an application of this method.
- *[See. R. Pease 2004 SC section]



TID Example

- Precision Voltage Readout for SDO Battery Charging Unit
 - Assess suitability of Linear Technologies RH1014 quad op amps
 - Issue: Very high impedance on the input makes application very sensitive to increased input leakage current—with parametric failure @ 15 nA
 - Application uses 10 parts, required $P_s \geq 99\%$ with 99% confidence

Parametric Failure Levels


mean failure level is 108 krad(Si)

standard deviation is 13 krad(Si)


lognormal mean=4.68

lognormal standard deviation=0.115

	Parametric Failure Level
Part 1	128.3 krad(Si)
Part 2	99.3 krad(Si)
Part 3	105.3 krad(Si)
Part 4	115.3 krad(Si)
Part 5	123 krad(Si)
Part 6	102.2 krad(Si)
Part 7	93.6 krad(Si)
Part 8	97 krad(Si)




Analysis using PCC and DMBP Method




Results are only good if the test sample is representative of flight parts.
Adding more data (e.g. more testing) would reduce the required margins.
More analysis (e.g. NOVICE TID calculation) could meet margin requirements.
Alternatively, we could mitigate by adding shielding or redesigning the circuit.

Similar methods apply to displacement damage.




SEE Requirements: Starting Point

- SEE requirements usually deal with allowable SEE rate
 - The more severe the consequences of an SEE, the lower the allowed rate
- Situation is different than for TID
 - Rate calculation methods use fit of σ vs. LET curve to a Weibull
 - Fits may be conservative or “tight”—yielding higher or lower rates respectively
 - Data for rare events (e.g. SELs or SEFIs) provide large error bars





SEL data for 256 Mbit SDRAM: WC fit rates are $\sim 5\times$ worse than best-fit rates




Poisson Errors

Errors on SEE cross sections scale as the inverse square of the counts on which they are based.




Systematic Errors

- Systematic errors need to be investigated and estimated. Failure of SiGe transistors to follow effective LET indicates the charge collection volume deviates systematically from the assumed RPP, and that systematic errors may be larger than normal.




After P. W. Marshall et al. [26]



Radiation Threat Evaluation

- For both degradation mechanisms and SEE, two issues important
 - How likely is the threat to be realized?
 - Measured as Probability of Failure, $P_f = 1 - P_s$ (P_s =Success Probability)
 - What is the impact if it does occur?
 - Severity is a qualitative measure (NASA System Engineering Handbook)
 - Category I—catastrophic
 - Category II—critical
 - Category III—Major
 - Category IV—minor
 - Failure Cost, C_f , is a quantitative measure of system impact
 - Usually monetary
 - Should include all costs
 - » loss of requirements, intangibles (loss of future business, etc.)
 - Specifying cost can be difficult
 - Note system impact influences what we consider acceptable P_f , and may even influence how conservatively we estimate it.
- We can combine probability and severity/cost in two ways
 - Criticality—qualitative specification of probability and consequences
 - Risk—quantitative, $R_f = P_f \times C_f$




Failure Risk and Failure Criticality


Criticality Table

P_f Severity	High Probability	Moderate Probability	Low Probability
Category I (Catastrophic)	Very Critical	Very Critical	Critical
Category II (Critical)	Critical	Critical	Moderately Critical
Category III (Major)	Moderately Critical	Moderately Critical	Acceptable
Category IV (Minor)	Moderately Critical	Acceptable	Acceptable


- Criticality—qualitative, simple, but hard to compare risks
- Risk—Complicated, but quantitative so we can direct/scale our effort to reduce risk
- Risk and Criticality are equivalent concepts
 - Assigning a cost to each criticality lets us use criticality as we do risk
- We phrase our analysis in terms of risk.



- System-level hardening becomes an exercise in risk reduction
 - Reduce R_f by reducing P_f or by reducing C_f or severity



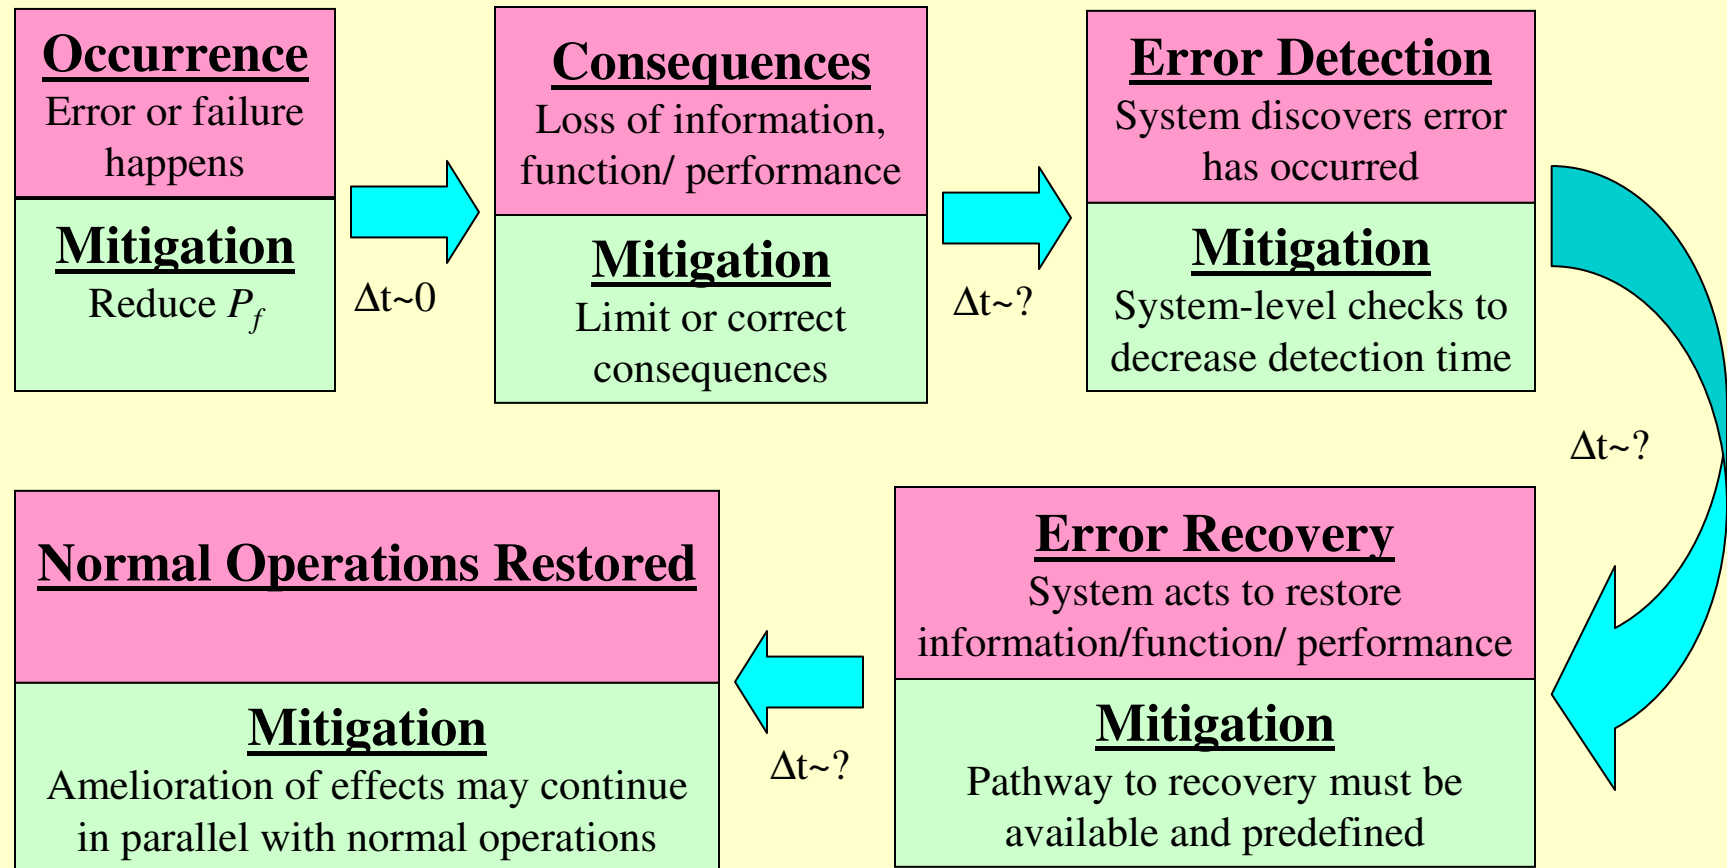
How Do We Decide We Need Hardening?




- Actual risk level unknown, so we calculate a bounding estimate
- Testing, analysis and mitigation lower risk, but at a cost
- Strategy driven by expected cost effectiveness of risk reduction



Stages of a Failure (and their mitigation)



In these 5 steps, only one of the mitigations involves reducing failure probability. The rest involve limiting consequences or hastening restoration of normal operations.




Radiation Effects and Their Consequences

Three Types of Effect

- Destructive SEE—SEL, SEB, SEGR, SEDR, SES, stuck bits
 - Depend mainly on technology
 - Result: WC—lose functionality of a single die; BC: partial functionality loss
 - System level: Worst-case—lose functionality; Best-case: reduce reliability
 - Effects limited to a single die; can happen any time
- Nondestructive SEE—SET, SEU, MCU, MBU, SEFI
 - Driven by cell function: Bi-stable cells—SEU, SEFI; Otherwise: SET
 - Results are application dependent:
 - SEFI: WC—functionality interrupted, data loss; BC—functionality interrupted
 - SEU, MCU, MBU: data corrupted
 - SET: WC—data corrupted; BC: No effect
 - Effects limited in space (SEFI, MBU, MCU, SEU) and time (SET)
- Degradation Mechanism—TID, Displacement Damage (DD)
 - Driven by both technology and application conditions
 - Result: WC: failure; BC: limited parametric degradation
 - System level: Worst-case—Application Dependent; Best-case: No effect
 - Effect is cumulative, but global—even affects unbiased “cold spares”

SEL: What Matters at the System Level


SEL Mechanism



After K. Galloway and G. Johnson [29]



- SEL is a regenerative, high-current, parasitic bipolar effect
 - SEL vulnerability increases with bipolar gain (and so with temperature)
- The substrate is important
 - High-energy ions yield higher SEL cross sections

- Interpreting SEL data is complicated
 - Part may have multiple modes, some destructive, some not
 - Testing needs to bound WC app. conditions: temperature, voltage, ion/proton energy, angle...
 - Latent damage must be investigated
 - Part must have multiple SELs
 - Microscopic examination
 - SEM and DPA
 - Post-SEL life test
- System-level effect
 - Renders a single-die inoperable
 - Latent damage renders parts unreliable
 - Nondestructive SEL is recoverable, but all data on part is usually lost.



SEGR and SEB: System Level Perspective

- SEGR occurs when increased electric field from ion-track holes pile up under the gate cause breakdown
 - MOSFETs vulnerable only when OFF
 - More vulnerable w/ high $|V_{DS}|$, $|V_{GS}|$
 - NMOS more vulnerable than PMOS
- Charge collection volume is not RPP
 - Cross section has complicated dependence on ion energy and angle
- SEGR testing difficult and expensive
- Renders single MOSFET inoperable.




After M. Allenspach et al. [44]

- SEB occurs in power BJT or MOSFET
 - Conditions similar to SEGR
- Can also render xstr inoperable




Others Destructive SEE: System Perspective





- SEDR—Dielectric breakdown in antifuses of one-time-programmable FPGAs renders a portion inoperable
 - Rates were low
 - Later generations less vulnerable
 - Mitigations exist
- Stuck bits—local deposition of dose renders a single bit unprogrammable
 - Rates to date have been low and
 - Annealing decreases stuck bit accumulation over time
 - Can be treated effectively as a permanent SEU
- SE Snapback—regenerative, bipolar parasitic effect in NMOS
 - may be an issue in hardened SOI
 - current limiting may help
- Bipolar failures
 - First seen for AD9048 in 1994
 - Subsequently seen in bipolar linear devices AMP01, OP??
 - Cross section highest at normal incidence
 - rates to date have been low
- Failures in FLASH Memories
 - Vulnerable during ERASE and WRITE operations
 - Operations involve high voltage due to charge pump, probably gate rupture
 - Rate is low enough that some devices could be used on orbit if WRITE and ERASE cycles limited.

Hardening for these is best done at the process level.



SETs: System-Level Perspective




SET from Q9 of National Semiconductor LM124 after reference 55.


- Disturbance of normal output due to SET is “short” and “local”
 - Effect depends on whether SET is captured by a bi-stable device downstream
 - Gives rise to strong frequency dependence
- SET susceptibility also depends on application conditions
 - bias, load, etc. all affect SET susceptibility



SEFI: The System-Level Perspective





- SEFI interrupt device functions and those of the system.
 - Usually SEFI result from errors in control logic, but...
 - ADCs with no control logic have also shown interrupts
- SEFI may also corrupt large amounts of data
 - WC SEFI corrupts all the data on a chip
- SEFI can be identified either by the loss of functionality large increases in data errors.
- Nondestructive SEL will look like a worst-case SEFI.



- Accumulation of errors in a 1 Gbit DDR SDRAM proceeds smoothly—shallow slopes indicating bit error, jumps are block errors. SEFI makes the error count literally go off the page.



SEU, MCU and MBU: System Perspective




After D. Radaelli et al. [64].


- Whether an SEE is a SEU, MCU or MBU depends on device interleaving
 - MCUs will look like multiple SEUs
- Consequences of SEU, MCU and MBU are all corrupted data
 - MBU require more sophisticated error correction.



Degradation: System Perspective




- TID application dependent
 - bias
 - load
 - temperature
 - etc.
- ELDRS
- May vary significantly lot-to-lot
 - sometimes even part-to-part
- Initial degradation negligible
 - significant at the system level only when application margins become exhausted
- Degradation is global.
 - entire part may be affected or only a portion of functionality
 - cold spares may degrade as quickly as primary parts




After Pease et al. [69].

- Micrel 29372 degraded more rapidly with zero load
 - also bias and temperature dependent




Degradation: System Perspective

- Displacement damage degrades performance due to
 - Minority carrier lifetime reduction
 - Reduced carrier mobility
 - Carrier removal
 - Increased leakage current
 - Thermal charge generation
 - Less variable lot-to-lot
 - Less application dependent
- Some parts susceptible to both TID and displacement damage
- Effect is global—affecting both primary and spare units
 - Hard failures rare but parametric degradation can be significant.




After Johnston et al. [73].

- Exception: Application dependent DD
- Annealing in amphoterically doped LEDs assisted by higher currents

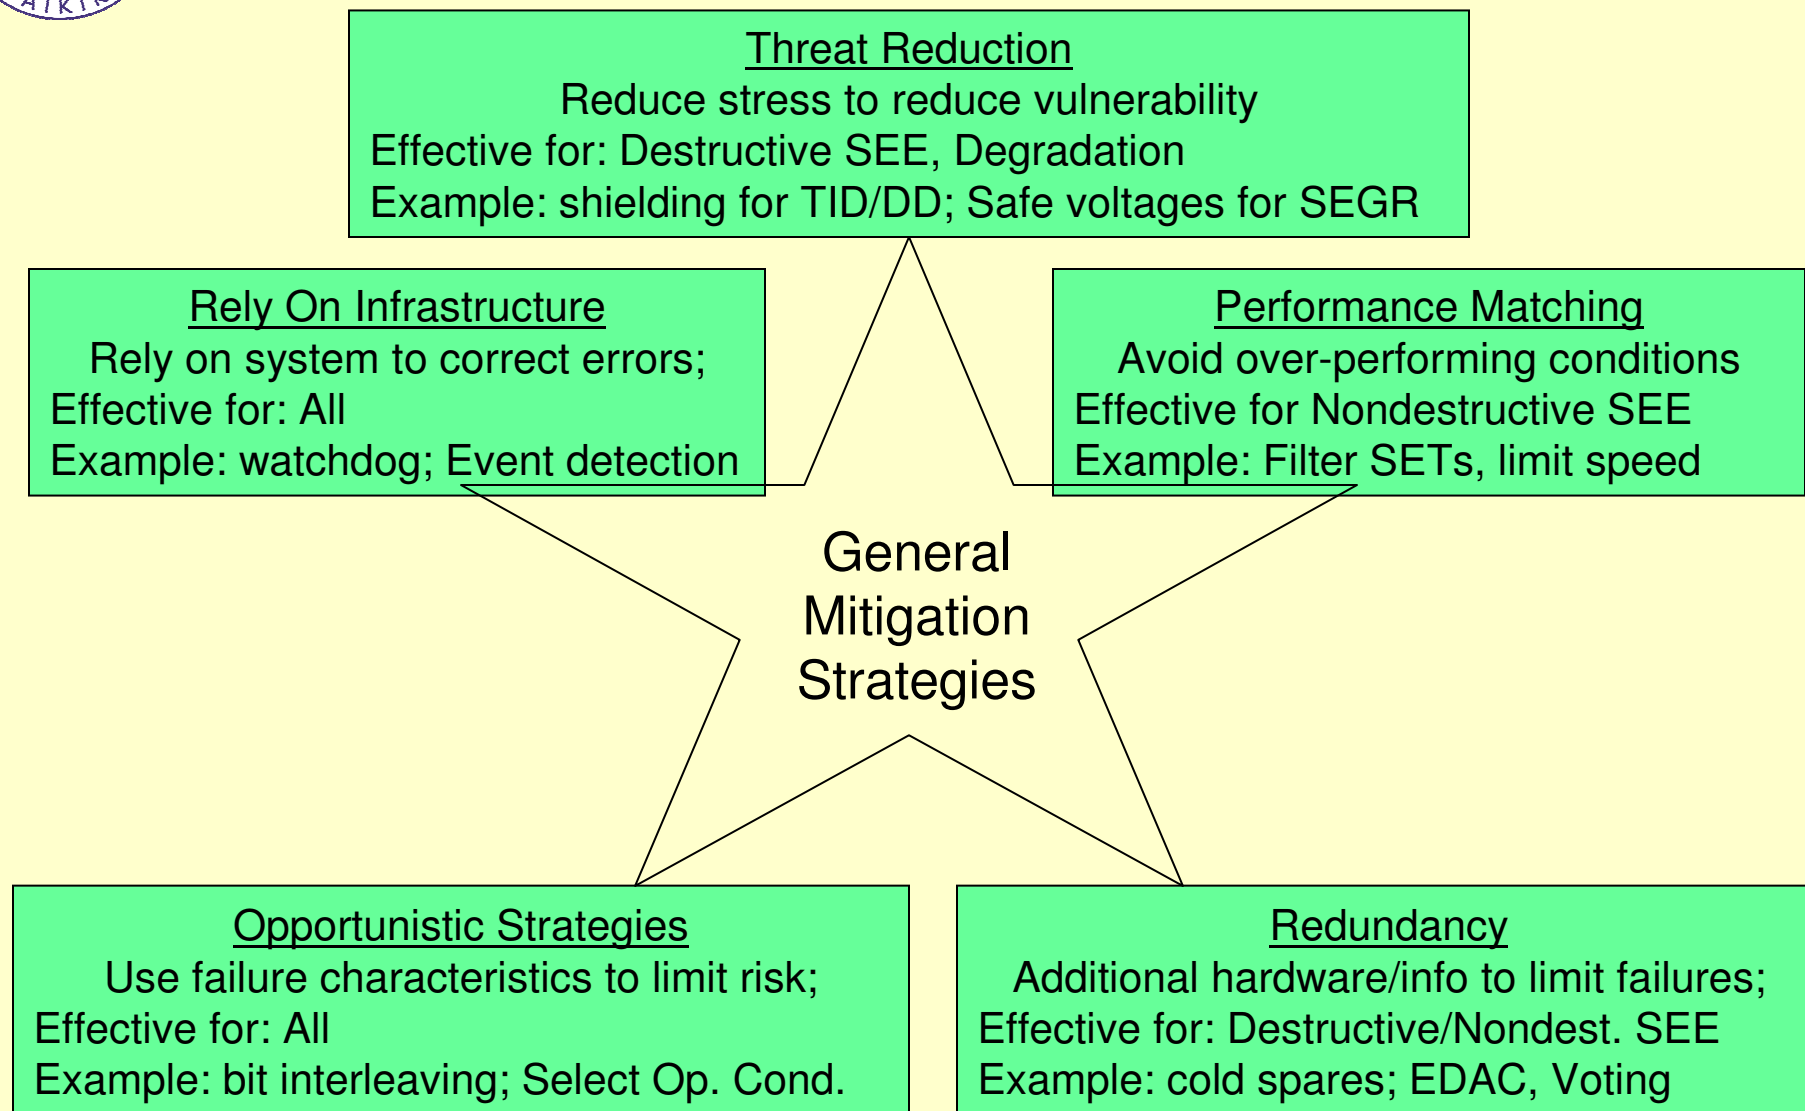



Radiation Effects at the System Level

- Radiation effects cause anomalous operation, resulting from
 - Abrupt failure of a device due to destructive SEE
 - Interruption of normal operations by a nondestructive SEE
 - Interruption may be at the part, circuit or system level
 - Data corruption/loss due to a nondestructive SEE
 - Degraded functionality due to TID or displacement damage
 - Failure (usually preceded by a period of degraded functionality) due to TID or displacement damage
- Mitigations of the above effects include efforts to
 - Keep the effect from happening (decrease its probability)
 - usually by changing environmental stress or application conditions
 - Speed up the discovery of the anomaly so recovery can begin
 - usually by having infrastructure in place to discover or speed recovery
 - Limit or ameliorate the consequences of the effect
 - May involve adjusting application conditions of repairing the damage



Types of Mitigation Strategies






Destructive Failures I

Threat Reduction

- SEGR and SEB: Only mitigation is operation at safe VDS and VGS
 - Empirically determined
 - For rad-hard parts rated $V_{DS} < 200$ V: For $V_{GS} \sim 0$ and $V_{DS} < 30\%$ of rated value, SEGR and SEB have not been seen
 - Caveat: Commercial devices have failed for $V_{DS} \sim 22\%$ of rated value
- Lowering operating temperature may reduce SEL susceptibility
- Shielding lowers TID and DD and decreases failure probability


Event Detection and Protection

- Used so far only with SEL
- Caveats
 - Must be effective against latent damage
 - Spurious SEL indications due to SET can lead to high outage rates.

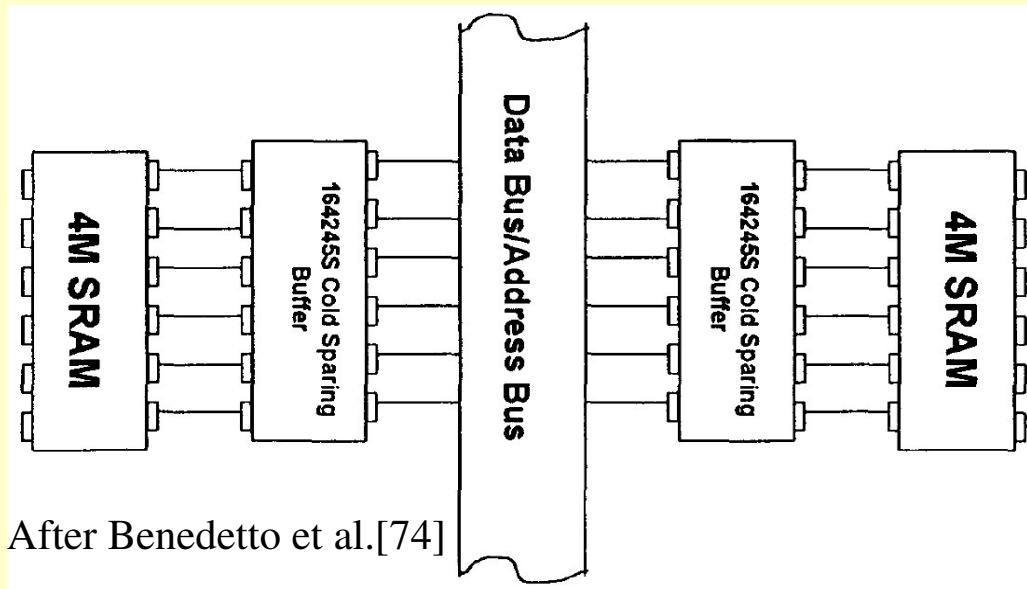





Destructive Failures II




- Cold sparing can extend mission life if hard failure risk in a critical system is unavoidable
 - Must isolate failed part and switch in redundant part
 - Full redundancy increases system reliability/life
 - Fully redundant means any part can replace any other
 - If expected life= T for nonredundant system, an $n:m$ fully redundant system should last $(n-m+1) \times T$
 - Need to make sure switching does not decrease total system reliability




After Benedetto et al.[74]

Cold sparing effectiveness is limited against TID or DD, since both primary and redundant parts degrade.




SEFI: Mitigating Loss of Functionality

- SEFI interrupt normal device and system functions
 - This doesn't mean that functions stop—they may malfunction
- Mitigation of SEFI
 - Most SEFI mitigations seek to identify a SEFI more rapidly
 - Most techniques are borrowed from fault-tolerant computing
 - Watchdog timer—forces a reset if system does not complete operations within a pre-allotted time
 - Error Counter—monitors the number of errors and forces a reset if the error rate exceeds a certain level
 - Health checks—requires the system interrupt operations and report on system health at regular intervals—e.g. heartbeat
 - Software can implement sophisticated monitoring
- Costs
 - Time spent monitoring is not spent on task
 - Some mitigations can also “upset”
- Mitigation of data loss due to SEFI is handled by separate means




Mitigating Data Loss: Voting

Most Data loss mitigation relies on redundancy: Voting is the classic example



Triplicate and Vote

- Triplicate and Vote yields an unambiguous answer
 - Reliable as long as common elements are hardened
- Cost in terms of power, space, weight is high
 - Cost is incurred even when no error occurs




Temporal Voting


- Much lower overhead (space and weight)
- Penalty is mainly speed
 - Third sample not needed if first two agree

Duplicate with Retry

- Compromise between overhead and speed
- Retry is necessary only when an error occurs



Voting schemes need not be majoritarian. Polling, averaging, etc. can all be viewed as voting systems that mitigate against different kinds of errors.




Voting Usually Works, But...

- Voting is the “big gun”—discussed in detail by Fernanda
 - effective, straightforward to implement, but requires high overhead
 - also, voting logic and other “domain-crossing” logic are still vulnerable
- Other strategies use Error Detection and Correction (EDAC)
 - implements redundancy as error-check bits
 - Much lower overhead, but can only correct a limited number of bits
 - Number of bits that can be corrected \leq half the number of error correction bits
 - Can be implemented to correct MBU, but probably not SEFI—at least by itself
 - EDAC can be supplemented with other techniques to handle SEFI
 - Problem is similar to burst errors in communications
- Which strategy?
 - Strategy choice is driven by requirements, cost, and system performance
 - If SEFI rate is low—EDAC may be sufficient by itself
 - For others, EDAC will need help
 - Sometimes voting is the only way to go



Hamming Code: An example of EDAC



$$E0 = D0 \oplus D1 \oplus D2 \quad (1)$$


$$E1 = D0 \oplus D1 \oplus D3 \quad (2)$$

$$E2 = D0 \oplus D2 \oplus D3 \quad (3)$$

Data bits				Check bits		
D0	D1	D2	D3	E0	E1	E2
0	0	0	0	0	0	0
0	0	0	1	0	1	1
0	0	1	0	1	0	1
0	0	1	1	1	1	0
0	1	0	0	1	1	0
0	1	0	1	1	0	1
0	1	1	0	0	1	1
0	1	1	1	0	0	0
1	0	0	0	1	1	1
1	0	0	1	1	0	0
1	0	1	0	0	1	0
1	0	1	1	0	0	1
1	1	0	0	0	0	1
1	1	0	1	0	1	0
1	1	1	0	1	0	0
1	1	1	1	1	1	1


- Adding 3 check sum bits to a 4-bit word gives us Hamming (7,4) code
 - 127 possible values, but only 15 are valid
 - Hamming distance=number of bits that change from a valid code to another
 - =3 for Hamming(7,4)
 - 1 bit flip—correct by going to nearest valid code
 - 2 bit flips—equidistant between 2 valid codes
- Hamming(7,4) is Single-Error-Correct-Double-Error-Detect (SECDED)


Bit in Error	Eq. (1)	Eq. (2)	Eq. (3)
D0	FALSE	FALSE	FALSE
D1	FALSE	FALSE	TRUE
D2	FALSE	TRUE	FALSE
D3	TRUE	FALSE	FALSE
E0	FALSE	TRUE	TRUE
E1	TRUE	FALSE	TRUE
E2	TRUE	TRUE	FALSE



Other EDAC Codes

- Many other EDAC codes are used for various applications.
- General characteristics
 - Codes can be implemented bit by bit or block by block (block=nibble, byte, word...)
 - Notation: (n,k) means the code has n bits or blocks, k for data, $n-k$ for EDAC
 - Efficiency= k/n
 - # of correctable bits $\sim (n-k)/2$ for k even $(n-k-1)/2$ for k odd
- Most codes used for satellites are generalizations of Hamming Codes
- Bose-Chaudhuri-Hocquenghem (BCH) codes represent data algebraically
 - Reed-Solomon codes are the most important example of BCH codes
- Data are encoded as symbols—blocks of nibbles, bytes, etc. for a Reed-Solomon code
 - Visualize the k data blocks $\{x_k\}$ as lying on a polynomial $P(x)$ of degree $k-1$
 - Error correction bits come from oversampling $P(x)$







Extending EDAC

- An EDAC code with k correction bits can correct up to $k/2$ bits in error
- But what if we have a SEFI that corrupts all the data on a single die?
 - Need to keep data loss on a single chip from overwhelming EDAC
- Answer: Interleave bits across die, just as we interleaved bits within a die to decrease MBU susceptibility

Example


- R-S(12 nibble, 8 nibble) can correct any 2 nibbles in error in any word
 - Store 1 nibble per word on each die
 - This means we can correct up to 2 worst-case SEFI, so if the SEFI rate is R , system rate $\sim R^3$
 - Could happen if the memory sits long enough
- Scrubbing means looking for errors and correcting them within time $T \ll 1/R^2$
- System can correct 2 WC SEFI (> than TMR)
 - Overhead is only 50%






Limitations of EDAC et al.

EDAC works great for memories—but less well for complicated devices




- For Processors, SEE corrupt algorithms as well as bits
 - internal error checks are also susceptible.
- For reprogrammable FPGA, the situation is even worse
 - Hardware itself can be changed.
- Interleaving is not possible, so EDAC has limited benefit.
- Visibility into errors is very limited
 - Many errors have no effect, while some cause the device to stop functioning
 - External error monitoring—watchdog timers, scanning, checksum, error counters, etc. are strongly recommended for such devices.
- Voting is more effective for devices executing complex algorithms. It can be done:
 - Internally—much lower overhead, but more likely to have common elements that will upset.
 - SET in combinatorial logic may give system errors a strong frequency dependence
 - Externally—high overhead, but offers greatest flexibility for hardening of voting circuitry.
 - It may be difficult to work out timing without impacting efficiency



Approaches to Mitigating Degradation

- Keep damage from happening
 - Shielding effective for TID in electron dominated environments
 - Less effective against protons and ineffective against Bremsstrahlung
 - DD usually proton dominated
- Minimize the damage
 - Ex.: limit damage in unbiased Micrel 29372s by alternating primary and redundant units

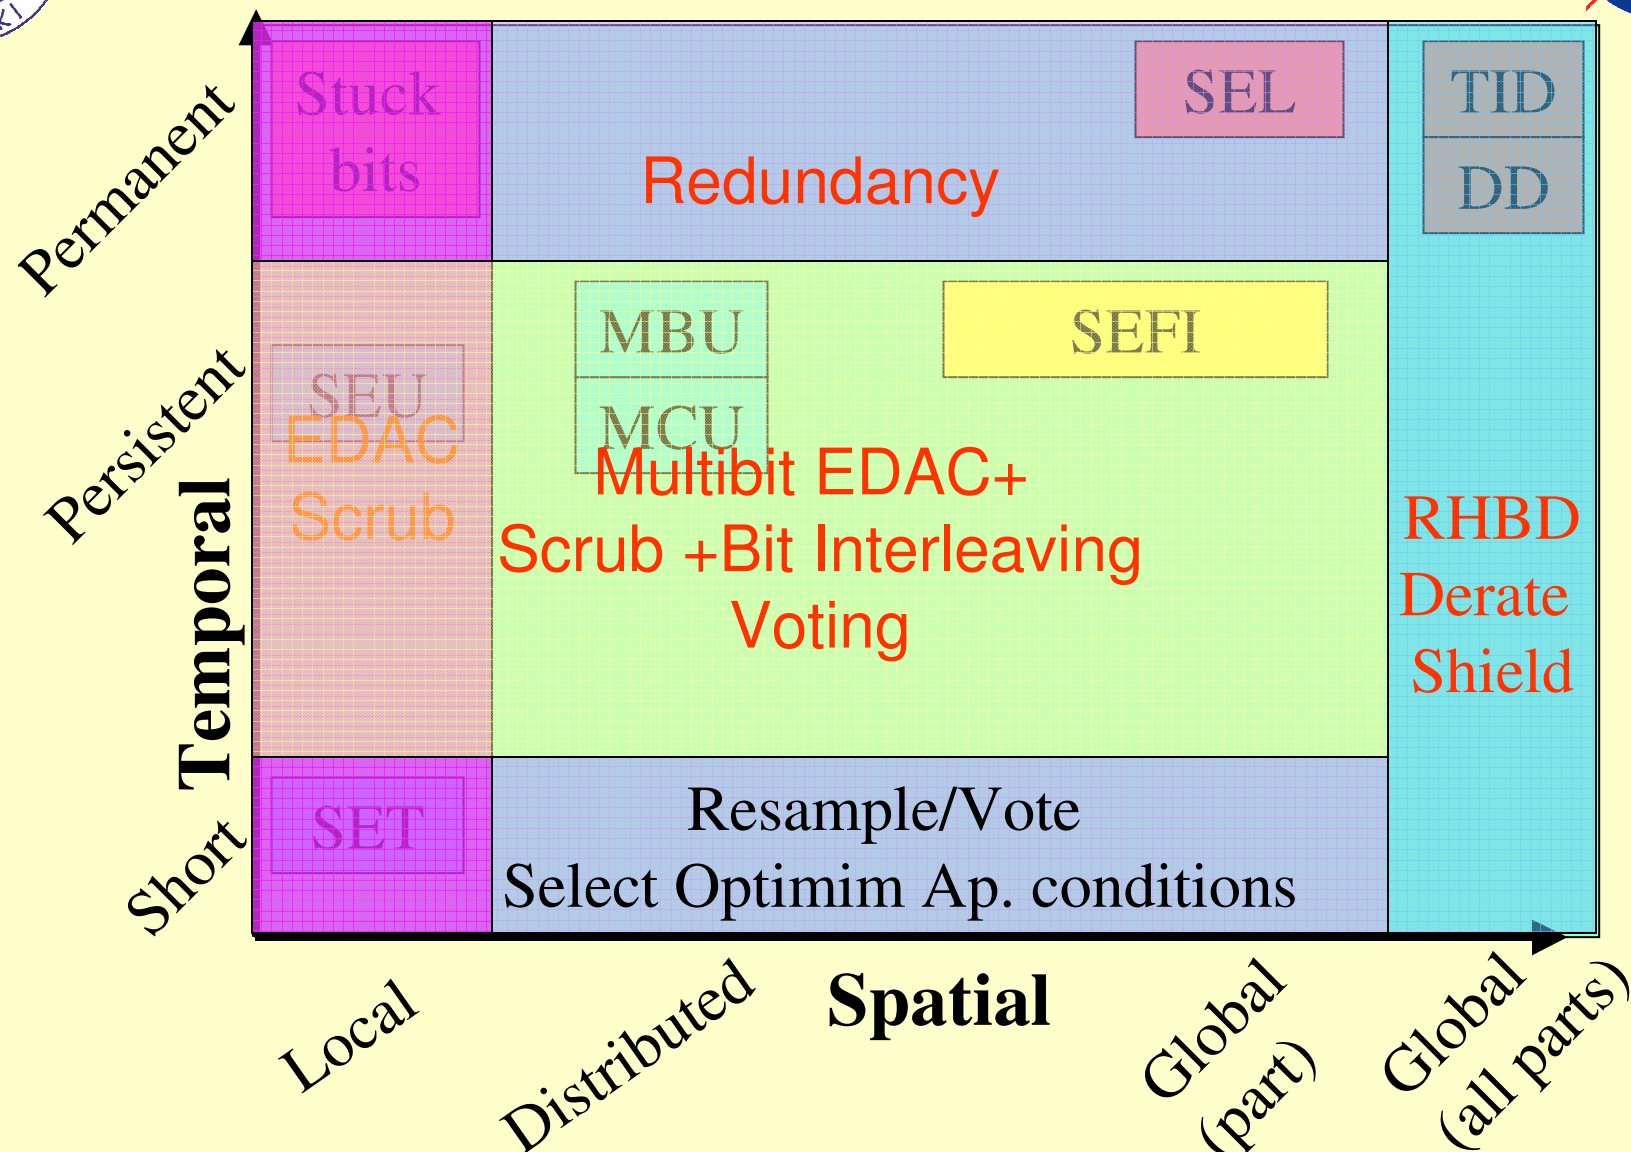
Dose-Depth Curve for GEO




- Accommodate damage
 - Ex: Increase design margins to ensure success at end of life
- Compensate for damage
 - Ex I: Include compensation circuitry to supply higher drive current at EOL
 - Ex II: Include ability to run at lower frequency to compensate for degraded timing




Combining Mitigations






Example: Hardening a Solid-State Recorder

- Requirements for Solar Dynamics Observatory SSR
 - 3 Gbit volatile memory and board space, power, weight... are limited
 - Must operate through solar particle events
- Memory size necessitates use of SDRAMs
 - Largest part available at the time was 256 Mbit SDRAM
 - Only one part from this generation had a reasonably low SEL rate
 - Hitachi/Elpida HM5225805B—×8 configuration preferred
 - Serious radiation issues
 - Large lot-to-lot TID variability—and lot traceability not available
 - Parts are susceptible to SEL, but only for $T > 50$ °C and no destructive SEL seen
 - SEFI, MBU, SEU and stuck bits occur at rates high enough to require mitigation
 - Questions:
 - How do we ensure lot traceability for TID hardness assurance?
 - Does the part have destructive as well as nondestructive SEL modes?
 - What about latent damage?
 - Do stuck bits accumulate enough to compromise EDAC at EOL?
 - Are SEFI rates low enough that we can operate through a solar particle event?
 - What mitigation scheme(s) are needed to ensure reliable operation



Resolutions: Procurement and Testing

- Highest priority for procurement was obtaining lot traceability
 - Worked with a value-added supplier (Maxwell Technologies), who bought a wafer lot of die and packaged them in their RadPak™
 - Ensures both lot traceability and shields parts to well above 2x margin
 - In RLAT one part failed at 40 krad(Si), with the passing above 50 krad(Si)
- Testing priorities: resolve issues with limited mitigation options
 - SEL: Part was latched >200 times by both heavy ion and laser
 - No evidence of destructive failure
 - Latent damage tests and analysis included 1000hr burn in, microscopic examination and DPA—no evidence of latent damage
 - For mitigation purposes: Nondestructive SEL looks like a WC SEFI
 - Stuck bits: Formation and annealing of stuck bits was examined
 - Formation rate was moderate, and most stuck bits annealed within minutes
 - Stuck bits are unlikely to be an issue
 - SEFI, SEU and MBU rates were also determined





4 Data Loss Mitigation Scenarios

- 1) No mitigation— ~10 outages in a 5 year mission
- 2) Reed-Solomon (12 nibble, 8 nibble), 8bits per die interleaved, full memory scrubbed daily— ~0.026 outages per mission
- 3) Same as 2), but only 4 bits per die interleaved— ~.0002 outages
- 4) Triplicate voting +daily scrubbing— ~0.021 outages per mission
- For Category III, scenario 2 is adequate.

Nondestructive SEL or WC SEFI interrupts system function and results in large data loss.


Severity is Category III, (major error).






Solar Event Rates

- Nominal SEE rates can go up a factor of 100-1000 during SPE
 - But the SPE heavy ion energy spectrum is soft
 - Shielding can reduce rates significantly






Validation




- Once we implement mitigation, we have to show it works
- Validation may be by test or by analysis
 - Types of mitigation most likely to require validation by test are those most difficult to model
 - SEL detection and protection and/or cold sparing
 - SET mitigation
 - Compensation circuitry for TID
 - Often mitigation is already based on the best test data available and additional testing would not improve confidence
 - Validation in such cases often based on modeling and analysis
 - Example: Analysis of RLAT TID data for SDO SDRAMs shows that dose levels will be less than half the 99/90 dose level.
 - Example II: Analysis shows scenario 2 mitigation reduces outages by >99.7%
 - Fault injection can be a very valuable tool
 - Useful for test planning as well as validation
 - May be essential for validating complicated systems w/multiple mitigation layers




When Mitigations Break Down

- Many of the mitigation strategies described predate the Space Era
 - Data loss mitigations are borrowed from Communications
 - Mitigations of hard failures borrowed from reliable design
 - SEFI mitigations borrowed from reliable computing
- Mitigation techniques are robust if based on accurate information
- Potential risks for System-Level Hardening
 1. Violations of assumptions underlying the mitigation
 2. Fidelity of test conditions to application conditions
 3. Discrepancies between test samples and mission components
 4. Synergistic effects that cause part behavior to change
 5. Cost challenges




Violations of Assumptions

- Pre 2003, $SET < 100 \mu s$
- Gave rise to rule of thumb
 - WC SET duration $< 100 \mu s$
 - ms-long transients[60],[43] violated this assumption
- Testing would reveal the issue
 - could increase mitigation
 - could replace part
- Rules of thumb risk surprises




After W. Heidergott et al.[86]

To be presented by Ray Ladbury at the NSREC Short Course, Honolulu, HI, July 23, 2007



See M. O'bryan et al.[43]




Violations of Assumptions

Rule of Thumb	Oops!
SET are recoverable	Transients >1.8 V could damage Actel RTAX-S FPGAs
rf devices are SET immune	Some SOTA CMOS can respond to ps-duration SETs
MOSFETs w/ $V_{DS} < 30\%$ of rated value immune to SEB/SEGR	IRF640 200 V commercial MOSFET fails due to SEB w $V_{DS} = 44$ V[43]
Bipolar ICs are immune to destructive SEE	Failures of AD9048[48], AMP01[50] and other bipolar parts[49]
CMOS device immune to ELDRS	Dose rate effects in CMOS[87]
CMOS devices immune to DD	Bulk damage in SDRAMs?[88],[89]
Indirect ionization is unimportant if threshold $LET > 15 \text{ MeVcm}^2/\text{mg}$	Scattering of W and other metals by light ions (Warren et al., TNS2005)
If a device is functional after SEL,	SEL can cause latent damage[38]
Others?	Stay tuned!


- Testing appropriate to the device and application is the only defense




Test Fidelity



- ELDRS is a classic example of lack of test fidelity to application
 - Degradation worsens for the LM124 even down to 1 mrad(Si)/s[66]
 - At high and low dose rates parts can degrade in opposite directions
 - Some parts fail at low dose rate, but not at high dose rate
- Other application conditions (esp. bias) also have significant effects
- Pre-irradiation Elevated Thermal Stress (PETS) shown to be




After Johnston et al. [66].




After McClure et al., REDW2000, p. 100




Test Fidelity



- Lack of test ions w/ GCR energies is another fidelity shortcoming
 - Fe ion range @ GCR peak > 14 cm
- GCR ions may cause MBU even on memories with bits interleaved
- Even interleaving bits across multiple die in a stacked memory is not enough, since ions could cause SEFI or MBU in multiple die



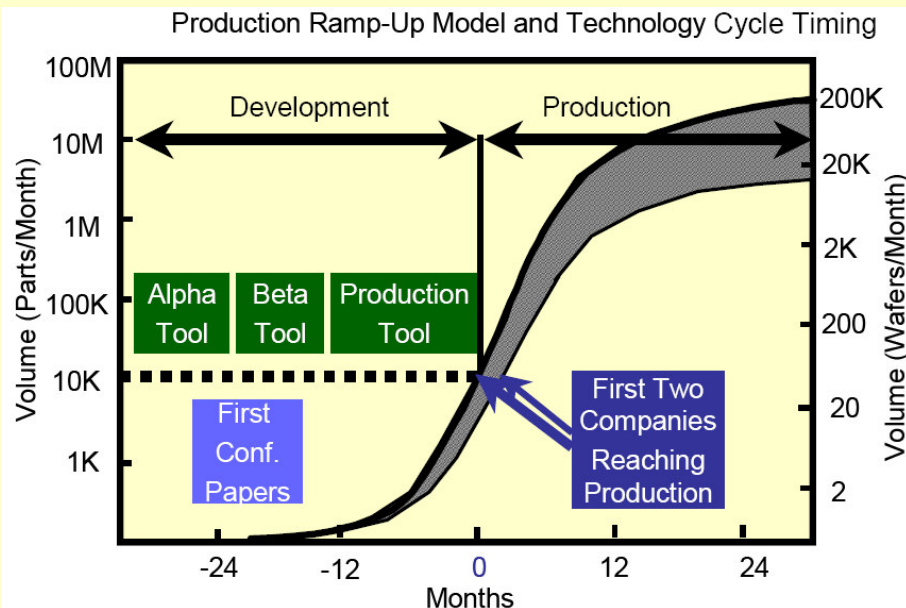
- Other test fidelity issues
 - Application dependence of SET
 - Mode dependence of SEFI, SEU in SDRAMs
- Modeling can help
 - But we need design information—supplied by manufacturer or reverse engineered
- Laser testing supplements heavy-ion testing to map out application dependence



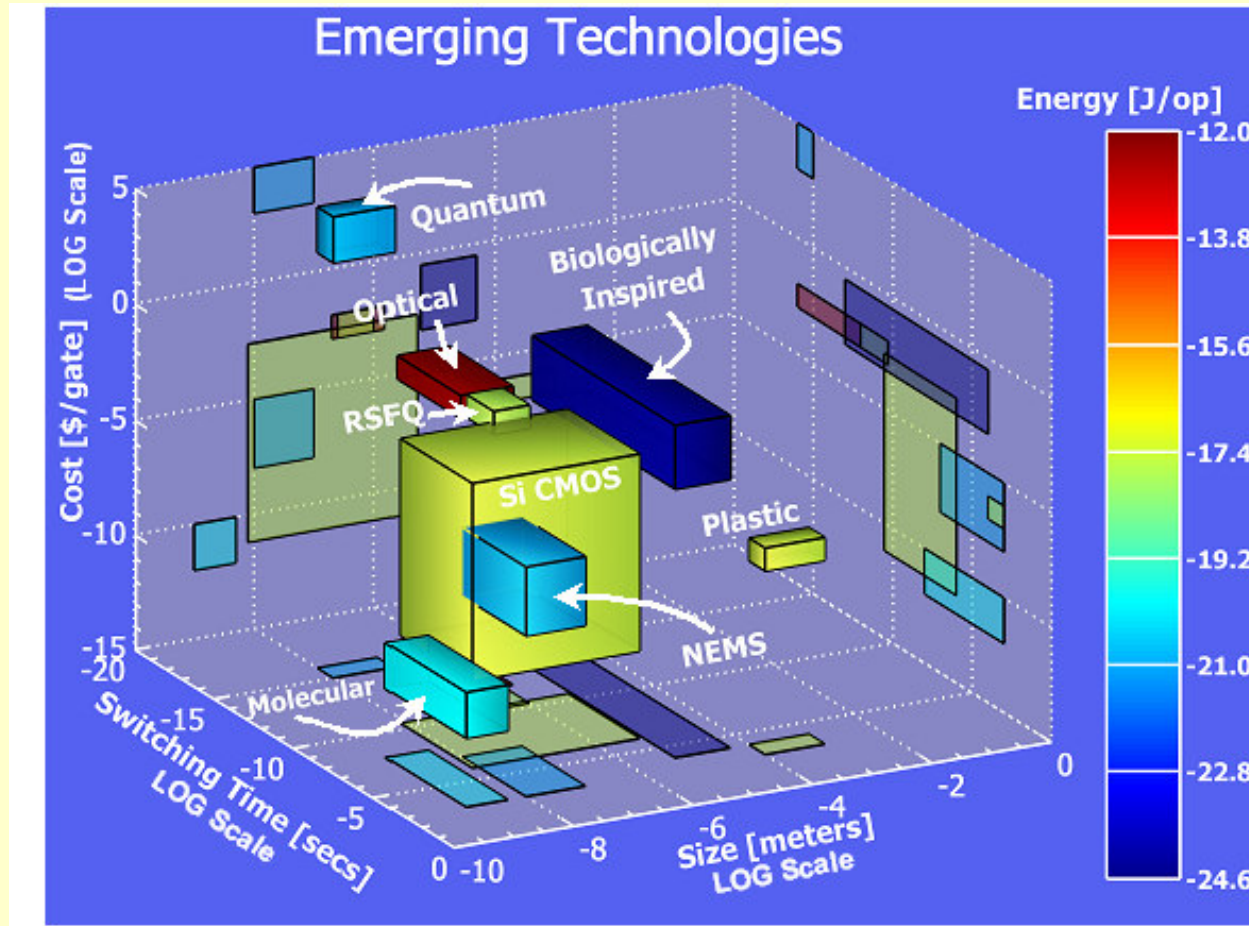
Ensuring Representative Test Samples

- Commercial technologies pose challenges beyond lot traceability
 - Short design and production cycle
 - Frequent die revisions
- Qualifying a commercial device takes time


- What if we qualify one revision only to find it is no longer available?
 - Requalify the new version and hope it meets our needs?
 - Look for the old die via other channels?
- Use of less formal procurement channels also poses risks
 - Electronics Resellers Association International identified 2857 independent brokers selling counterfeit parts
 - May have correct markings, pass initial screenings and even have fake certificates of compliance
 - May be recycled off of old boards
 - May be rejects with new markings
- For more information, see D. Meshel (Aerospace Corp).



What About The Future?



- One thing we have going for us—the devices have to work reliably to begin with
 - Radiation responses will be failure, interrupted functions, corrupted info or degradation
 - Questions include degree of variability and how much mitigation is needed.




Synergistic Effects

- Question is whether other factors change radiation response sufficiently to undermine mitigations based on tests of new devices
 - Aging can affect TID response, but effects so far have not been large
 - Humidity can affect testing for non-hermetic parts
 - Pre-irradiation Elevated Thermal Stress has been shown to have important effects on ELDRS response of some parts.
 - Other issues:
 - Do TID or Displacement Damage influence SEE response
 - See paper PJ-2 by Buchner et al. for an example
 - Does TID influence Susceptibility to stuck bits in DRAM
- Study of synergistic effects is still new
- Important issue is making sure testing includes these effects so mitigation strategies are effective both Beginning and End of Life



Keeping System Hardening Economical I



1996 SEE Test of a 4M SRAM				
Description	Man-weeks or units	Cost in \$	Total	Note
Heavy Ion at BNL SEUTF				
Test plan	0.20	\$4,000.00	\$800.00	Includes eng, rad, other to define what needs to go into test set with project.
Device procurements	10.00	\$50.00	\$500.00	
Misc parts	1.00	\$250.00	\$250.00	Sockets, connectors, etc...
Device delidding	0.05	\$3,500.00	\$175.00	
Test board design - electrical and layout	0.40	\$4,000.00	\$1,600.00	
Board fab and population	1.00	\$3,500.00	\$3,500.00	In-house board build
Board/tester debug	0.50	\$4,000.00	\$2,000.00	
Rad expert (test oversight and plan)	0.40	\$5,000.00	\$2,000.00	
Heavy ion test performance - contractor	2.00	\$1,500.00	\$3,000.00	
BNL Beam	6.00	\$700.00	\$4,200.00	Simple data: bit flips, latchup
Data analysis	1.00	\$3,500.00	\$3,500.00	
Test report (eng, rad expert, rad lead)	0.50	\$4,000.00	\$2,000.00	
			Total:	\$23,525.00


1996 vs 2006 a >3X Cost Delta

After LaBel, 2007[65]


2006 SEE Test of SDRAM				
Description	Man-weeks or units	Cost in \$	Total	Note
Heavy Ion at TAMU				
Test plan	1.00	\$4,000.00	\$4,000.00	Includes eng, rad, other to define what needs to go into test set with project.
Device procurements	10.00	\$75.00	\$750.00	
Misc parts	1.00	\$1,000.00	\$1,000.00	Higher speed drives cost
Device thinning and package processing	10.00	\$500.00	\$5,000.00	Assumes FBGA package; If this does not work, more expensive test facility like NSCL needed: >\$100K delta
Daughterboard Board design - electrical	0.80	\$4,000.00	\$3,200.00	
Daughterboard Board design - PCB	0.80	\$3,500.00	\$2,800.00	
Test Boards	10.00	\$500.00	\$5,000.00	
Board population	0.40	\$3,500.00	\$1,400.00	
Board/tester debug	0.50	\$4,000.00	\$2,000.00	
Tester VHDL development	4.00	\$4,000.00	\$16,000.00	
Technician	1.00	\$3,500.00	\$3,500.00	
Rad expert (test oversight and plan)	0.60	\$5,000.00	\$3,000.00	
Heavy ion test performance - contractor	3.00	\$2,000.00	\$6,000.00	
TAMU	16.00	\$750.00	\$12,000.00	2X time required: more data, more error types, more complex results; partial test
Data analysis	3.00	\$3,500.00	\$10,500.00	
Test report (eng, rad expert, rad lead)	1.00	\$4,000.00	\$4,000.00	
			Total in \$	\$80,150.00



Keeping System Hardening Economical II




- Increasing test costs are a concern because
 - They mean that we will probably have to do less testing
 - They can stifle innovation in design
 - They may delay projects consulting radiation experts
- System-level hardening helps contain costs by focusing test efforts around system effects—failure, SEFI, data corruption and degradation
- Other cost-containment strategies
 - Early involvement
 - Eliminate unnecessary risks, start RHBD efforts, map out system hardening
 - Cooperate for strategic needs across organizations
 - If the cooperation builds trust, it may facilitate data sharing as well
 - Develop innovative testing and qualification strategies
 - Example: Supplement SEE testing using Laser testing over application conditions
 - Integrate testing and device modeling
 - Testing feeds data into the models and validates its results
 - Modeling extends test results to conditions of application fidelity




Conclusions I

- There are no really new system-level hardening strategies because
 - System hardening concentrates on mitigating system-level effects
 - Hard failure—strategies borrow from techniques of reliable design
 - Functional interruption—strategies borrow from fault-tolerant computing
 - Data corruption—strategies borrow from communications
 - Degraded functionality—strategies borrow from techniques of reliable design
 - The techniques we have work and are applied in new, creative ways.
 - System hardening can achieve a harder system than rad hard parts alone
- Techniques fall into five broad categories
 - Threat reduction—keep the threat from happening by reducing stresses
 - Performance matching—avoid overperforming where threat is more likely
 - Redundancy—minimize consequences with redundant function or info
 - Opportunism—use threat characteristics to minimize its consequences
 - Infrastructure—accept that threat will happen and rely on infrastructure to speed recovery or minimize consequences



Conclusions II

- But system hardening is expensive in \$\$ and system performance
 - Understand and harden to requirements
 - Understand the threat to those requirements
 - Choose techniques that reduce risk most cost effectively
 - Validate the effectiveness of the mitigation
- System hardening has a future as long as there are space systems
 - New systems will have to meet new requirements
 - Radiation environments will continue to threaten those requirements
 - New technologies will pose new threats, but at the system level
 - Effects will still be: failure, interrupted service, data corruption and degradation
 - Mitigations will rely on the same strategies—hopefully improved
- Personally, I can't wait to see the strategies they use to fly a molecular memory



Detail of System-Level Hardening

